

Coupled 802.11 Flows in Urban Channels: Model and Experimental Evaluation

Joseph Camp, *Member, IEEE*, Ehsan Aryafar, *Member, IEEE*, and Edward Knightly, *Fellow, IEEE*

Abstract—Contending flows in multihop 802.11 wireless networks compete with two fundamental asymmetries: 1) channel asymmetry, in which one flow has a stronger signal, potentially yielding physical layer capture; and 2) topological asymmetry, in which one flow has increased channel state information, potentially yielding an advantage in winning access to the channel. Prior work has considered these asymmetries independently with a highly simplified view of the other. However, in this paper, we perform thousands of measurements on coupled flows in urban environments and build a simple yet accurate model that *jointly* considers information and channel asymmetries. We show that if these two asymmetries are not considered jointly, throughput predictions of even two coupled flows are vastly distorted from reality when traffic characteristics are only slightly altered (e.g., changes to modulation rate, packet size, or access mechanism). These performance modes are sensitive not only to small changes in system properties, but also small-scale link fluctuations that are common in an urban mesh network. We analyze all possible capture relationships for two-flow subtopologies and show that capture of the *reverse* traffic can allow a previously starving flow to compete fairly. Finally, we show how to extend and apply the model in domains such as modulation rate adaptation and understanding the interaction of control and data traffic.

Index Terms—Access protocols, ad hoc networks, analytical models, channel asymmetry, information asymmetry, measurement, physical layer capture, wireless mesh networks.

I. INTRODUCTION

IN URBAN environments, IEEE 802.11 nodes interact in many ways, e.g., within and among paths in a multihop network and among deployments from different domains. Competing transmitters rarely have equal link quality to a given receiver, i.e., *channel asymmetries* are prevalent, especially in urban channels. When packets overlap in time, even slight link quality differences have been shown to cause physical layer capture such that the packet sent over the higher-quality link is received correctly but the packet sent over the weaker link is dropped [1]. Moreover, transmitters or receivers of competing

flows often have unequal channel state information, a situation termed *information asymmetry*. In such cases, a topological asymmetry results in a hidden node having inferior channel availability information (see definition below and Fig. 5), forcing the hidden node to contend at random times guided by binary exponential backoff rather than at “idle times” driven by carrier sense. However, while the effects of information asymmetry and channel asymmetries are understood in isolation ([2]–[5] and [6]–[8], respectively), their interdependencies have thus far been ignored.

In this paper, we *jointly* consider information and channel asymmetries with both analytical models and extensive, urban measurements. As in [9], we employ the two-flow enumeration technique of [4] and consider all topological couplings of paired flows. However, in contrast to [4], [9], we inform the model of channel asymmetries via a signal-strength matrix between nodes and a measurement characterization of physical-layer capture events. By doing so, we reveal complex interdependencies of different system parameters that have been ignored by prior work. As an example, we show that not only do traffic parameters such as modulation rate and packet size change the timing of the model (e.g., backoff, carrier sense, and other factors that are most affected by information asymmetries), but they can also change the ability of traffic to perform physical layer capture at certain channel-asymmetry states. As a result, a new dimension emerges for throughput sharing that must be understood when attempting to achieve fairness in multihop wireless networks, both for planning and ongoing management of deployments.

In particular, our contributions are twofold. First, we develop an analytical model that predicts the throughput of coupled flows with the input of signal strengths between nodes. Using an embedded Markov chain to characterize a broad set of link-interaction states, we incorporate key system factors such as topology, modulation rate, packet size, channel conditions, and physical layer capture. This model is the first to jointly consider information and channel asymmetries when predicting the throughput of 802.11 flows. By doing so, we characterize our experimental finding that even with high-quality links, small-scale channel fluctuations common to urban environments can flip the throughput-sharing modes of coupled flows. One such example occurs when flows compete asymmetrically due to topological connectivity factors and unequal *channel availability information*. Namely, as shown in [10] and modeled with idealized channels in [4], a flow can starve due to lack of knowledge at the sender about when to begin contention. The flow with full information “wins” the contention nearly all the time. However, we show that by incorporating channel asymmetries, a favorable channel state at the receiver for the information-poor transmitter can allow that flow to receive equal throughput compared to the information-rich transmitter.

Manuscript received January 29, 2011; revised October 16, 2011; accepted November 24, 2011; approved by IEEE/ACM TRANSACTIONS ON NETWORKING Editor K. Papagiannaki. Date of publication January 23, 2012; date of current version October 11, 2012. This work was supported by the NSF under Grants CNS-0751173, CNS-08102501, and CNS-0619797; the Cisco Collaborative Research Initiative; and the Texas Instruments Leadership University Program.

J. Camp is with the Electrical Engineering Department, Southern Methodist University, Dallas, TX 75275-0338 USA (e-mail: camp@smu.edu).

E. Aryafar is with the Department of Electrical Engineering, Princeton University, Princeton, NJ 08544 USA (e-mail: earyafar@princeton.edu).

E. Knightly is with the Electrical and Computer Engineering Department, Rice University, Houston, TX 77005 USA (e-mail: knightly@rice.edu).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TNET.2011.2181863

Conversely, we show that with flows previously assumed to have equal throughput distributions due to symmetric information, only slight channel asymmetries cause one flow to achieve nearly zero throughput.

Second, we design a set of urban experiments consisting of thousands of measurements on a deployed urban mesh network. We first validate the analytical model and show that it is accurate in predicting the throughput of coupled flows for diverse channel conditions and topologies. Furthermore, our measurements indicate that the throughput sharing of many coupled flows vary widely over time. By examining the channel conditions associated with the maximum and minimum differences in flow throughputs, we empirically identify the small-scale channel fluctuations that cause such changes in throughput-sharing modes. Throughout, we show that *reverse traffic* (acknowledgment and clear-to-send packets traveling in the reverse direction of data) has a critical impact that has not been studied. In contrast to the data direction, this reverse channel is *not* carrier-sensed before transmitting. Thus, even within carrier-sense range, capture relationships have a critical impact yielding new link interdependencies, interactions with forward traffic, and vulnerable subtopologies, all characterized by the model. Finally, we discuss extensions to the model and how to apply it to the two domains of modulation rate adaptation and predicting the effect of control traffic on data flows.

This paper is organized as follows. First, we measure link variation and capture behavior of our large-scale urban mesh network in Section II to inform our model. We then present our analytical model in Section III. In Section IV, we perform an extensive set of experiments on an urban mesh network to both validate and apply our model. We then compare to related work in Section VI. Lastly, we conclude in Section VII.

II. BACKGROUND: TFA LINK VARIATION AND PHYSICAL-LAYER CAPTURE BEHAVIOR

In this section, we explore physical layer capture in an urban mesh deployment, the Technology For All (TFA) network. TFA¹ is a large-scale urban mesh network deployment that covers thousands of users spanning multiple square kilometers. At the time of the study, there are 17 access points that form a backhaul tier to provide Internet access to the aforementioned coverage area. Moreover, in a controlled, in-lab environment, we perform experiments with the same wireless card used in the deployment to understand relative link quality differences that lead to physical layer capture as a function of modulation rate and packet size.

A. Background: Timing Impact on Capture

Prior work has shown that the timing of the competing packets plays a critical role in physical layer capture due to the Message-in-Message (MIM) function required by the 802.11a standard and implemented in the Atheros chipset [7]. Namely, if a packet's preamble is received correctly (or, more precisely, enough synchronization bits within the preamble are received correctly), the receiver "locks on" to that packet (Fig. 1) and only switches to attempt to decode a later overlapping packet if that packet is greater than 10 dB stronger than the first transmission. As a result, there are two capture thresholds based




Fig. 1. Node A leads B by the synchronization bits in the preamble, which allows the receiver to lock on to A's packet.




Fig. 2. Capture experiment setup for Prism 2.5 Chipset.

upon timing. However, in our experiments, we use the Prism 2.5 chipset, which has been shown to not implement MIM and forces overlapping transmissions to result in loss if the receiver is locked on to the weaker packet [1]. The ability of the sender to capture depends completely upon the correct Frame Check Sequence versus an interfering source. Thus, in the case where the capture occurs for the Prism chipset (i.e., the stronger packet is first or the stronger packet trails the weaker packet by less than the synchronization bits of the preamble, lasting six slots in 802.11b), a single capture threshold exists for the physical layer rate and packet size combination. Our analysis in this paper models this family of capture behaviors. However, the same methodology can be extended to include MIM.

B. Capture Experiment Setup of Prism Chipset

To inform our model about the physical-layer capture behavior in TFA, we measure the delivery ratio of the wireless card when competing against another transmitter at certain relative signal-to-noise ratio (SNR) values. The wireless card is the SMC EliteConnect 2532-B, which is an 802.11b card using the Prism 2.5 chipset. Fig. 2 depicts our use of both channels of a Spirent SR5500 channel emulator for the transmissions of the node trying to be captured (sender on channel 1) and the competing transmitter (interferer on channel 2). The sender and interferer are unable to carrier sense one another. For the interferer, we eliminate the effects of binary exponential backoff by sending infinitely long packets. For the sender, we use broadcast traffic to not need ACKs (since we are unable to create a third channel in the reverse direction). We combine the two channel outputs and connect them to the input of the receiver.

For each relative value of SNR, we hold the power constant for channel 1 (-72 dBm) and vary the power of channel 2 (-70 to -84 dBm), testing a relative SNR range from -2 to 12 dB. During the measurement, the relative power levels are held constant (i.e., no fading is experienced on the channels). We send equally spaced broadcast packets (nearly flooding the channel) over a 60-s duration, running synchronized scripts which record the card statistics for both sender and receiver before and after the experiment. Thus, we are able to calculate the packet delivery ratio as the total amount of received packets over the total transmitted packets.

C. Capture Threshold Per Modulation Rate and Packet Size

Prior measurement studies on physical layer capture have shown that higher modulation rates require greater relative SNR to achieve capture [7], [8]. Since TFA has a diverse traffic profile

¹For a description of the network and hardware, refer to <http://tfa.rice.edu>.




Fig. 3. Capture probability of TFA hardware in lab on two controlled channels of a channel emulator. (a) 1500-B packets. (b) 2 Mb/s.

with large-size, high-rate packets interacting with small-size, low-rate packets we must experimentally understand the role of both modulation rate *and* packet size, which has not been fully explored. To choose relevant packet sizes, we refer to recent studies on the Internet that have classified the traffic according to three different groups of approximately 100, 500, and 1500 B [12]. We additionally consider the case of 1000 B. We measure the capture thresholds for all four physical layer rate and packet size combinations (16 configurations). These measurements inform our model about the physical layer behavior based upon both factors.

Fig. 3 depicts the delivery ratio for the sender at each modulation rate for 1500-B packets (left) and each packet size for the control rate (right) according to each relative SNR value from -2 to $+12$. For 1500 B (a common TFA packet size), observe that nearly zero packets are delivered for 0 dB and below (when the interferer is as loud or louder than the sender). As the difference in SNR increases, the lowest modulation rate quickly converges to nearly 1 for 5 dB, yet the highest modulation rate is only able to obtain close to 1 for a relative SNR of 12 dB. In contrast, for the control rate in the right figure (2 Mb/s), 2 dB is enough to capture 80% of the packets for the smallest size (100 B). A difference of over 6 dB is required for the same performance of the largest-sized packets (1500 B). Referring back to the 1500-B result, the relative SNR required to achieve the same packet delivery ratio between a small-sized, low-rate frame and a large-sized, high-rate frame is up to 8 dB different. Thus, the ability to capture is highly dependent upon both packet size and modulation rate of the strongest overlapping packet.

D. Capture Prevalence in TFA

Based on the in-lab measurements for when capture occurs, we now consider capture relationships across the network based on a distribution of relative signal quality of competing link pairs along the TFA backhaul tier. We consider a single point in time, though we have verified that similar distributions exist for all of our measurements that span a week's time with per-second measurements over 10-min durations. At the beginning of each test interval, the mesh nodes are synchronized according to a global clock using *ntptime*. From the signal measurements, we sequentially search each mesh node for any two possible links that would compete at a mutual receiver. Since reverse traffic does not carrier sense and since these cards have been shown in [13] to lack the function of physical carrier sensing, we do not exclude competing in-range link pairs from consideration.

Fig. 4 shows the number of competing link pairs in the TFA backhaul tier according to relative-SNR groupings. At the time of the measurement, there are 1621 total competing link pairs




Fig. 4. Histogram of all competing backhaul link pairs within TFA according to relative SNR.




Fig. 5. Snapshot of coupled flows with (a) symmetric and (b) asymmetric cross-flow connectivity where heterogeneous modulation rates are used per flow.

from 17 omnidirectional access points (each also serving as a receiver). Of these link pairs, over 72% have relative SNR differences of greater than 2 dB, resulting in 97% capture for control traffic. Thus, there is a high degree of physical layer capture that occurs in the TFA topology.

III. MODEL

In this section, we develop a bidimensional, discrete-time Markov chain embedded over continuous time to study the throughput sharing behavior of two coupled flows. We explicitly account for different capture relationships that exist among competing nodes. The system state is the joint backoff stage of the coupled sources. The transition probability is determined by capture relationships and other system parameters, allowing different capture relationships to drive the steady-state distribution of the system. Using our model, we accurately predict the throughput and investigate the impact of capture relationships and other parameters on the system performance.

A. Background

We study the performance of coupled flows in multihop wireless networks. Fig. 5 depicts a snapshot in time of such flows with symmetric and asymmetric cross-flow connectivity.

Coupled and Uncoupled Flows and Hidden Terminals. In most cases, a flow such as Aa that is interfering with flow Bb has backoff behavior that is *coupled* to that of flow Bb . In other cases, such as with broadcast traffic, the backoff processes of competing flows are *uncoupled*. When two transmitters such as A and B use 802.11, if *intersender* interference exists in which packets can be sensed or decoded between transmitters, one transmitter defers while the other transmits. The resulting behavior can be predicted using existing models such as [14] and extensions. However, where the two transmitters have no inter-sender interference (hidden terminals), no prior model jointly considers channel and information asymmetries.

Cross-Flow Connectivity: Flows with coupled backoff behavior can have cross-flow connectivity where the sender of one flow receives packets from the receiver of the competing flow. There can be *symmetric* [Fig. 5(a)] or *asymmetric* [Fig. 5(b)] *cross-flow connectivity* if the senders of both flows can decode packets from the receiver of the competing flow or if only one sender is able to do so, respectively. The symmetry or asymmetry of this relationship has been shown to cause balanced or imbalanced throughput sharing with idealized channels due to the MAC layer behavior [4], [10] and equal or unequal channel availability information, respectively.

Complexity of Capture Relationships: Node mobility or environmental movement can cause fluctuations in channel quality. Furthermore, there are spatial differences between the senders of competing flows from a given receiver. The resulting difference in link qualities can cause physical layer capture for competing transmitters that are out-of-range and, thus, can have simultaneous transmissions [1], [7].

Physical layer capture can occur for traffic in the forward direction (e.g., data or RTS packets from A and B overlapping at a in Fig. 5) or for traffic in the reverse direction (e.g., CTS or ACK packets from a and b overlapping at A in Fig. 5). For a given flow, there can be *forward traffic capture* over the forward or reverse competing traffic and *reverse traffic capture* over the forward or reverse competing traffic. There are a total of four possible capture scenarios for a given flow with respect to competing flows and three possible results: winning capture, losing capture, or collision with loss. To find the capture relationships per topology, we use the received SNR at each one of the nodes combined with the empirically derived capture behavior of the card (as discussed in Section II-C). One could approximately predict such an SNR matrix by using the transmitter characteristics (e.g., transmit power and antenna properties) with well-known path loss models [15].

B. Analytical Model

Joint Channel State Evolution: In order to correctly analyze the behavior of coupled sources, we consider the joint backoff evolution of the two flows. We identify three main states: 1) an idle state (σ); 2) a single-access state (T_S); and 3) an overlapping-packets state (T_θ). In the single-access state, either one flow transmits or both flows transmit, but the first packet of the earlier flow does not overlap with other flow's later packet (e.g., in Fig. 5(b), A 's RTS finishes and is receiving a CTS while B starts transmitting an RTS). In the overlapping-packets state, the first packets of both flows overlap in time. While σ is a constant equal to one mini-slot in 802.11, the duration of the other intervals (T_S and T_θ) depends on the modulation rates of transmitting nodes, the access mechanism, and the duration of overlapping packets. For example with RTS/CTS enabled, the duration of a successful single-access state is equal to

$$T_{S_n} = \frac{\text{RTS} + \text{CTS} + \text{ACK}}{R_{\text{basic}}} + 3 \cdot \text{SIFS} + \text{DIFS} + \frac{l_n}{R_n}. \quad (1)$$

In the above equation, the packet size and modulation rate of node n are denoted by l_n and R_n , respectively. In the overlapping-packets state, one or both packets are captured or both packets are dropped. Thus, the state duration is highly dependent upon capture relationships as well as other system parameters. These values are computed for each case once their corresponding probabilities are calculated.

TABLE I
BASIC ACCESS WITH SYMMETRIC CONNECTIVITY

To State	Transition Probability
i, j	$(1 - \gamma_i)(1 - \gamma_j)$
$0, j$	$\gamma_i(1 - \gamma_j)^{f_A}(1 - l_A)$
$0, j + 1$	$P_\theta S_{AB} C_{ABa,D}(1 - l_A)$
$i + 1, j + 1$	$P_\theta S_{AB}(1 - C_{ABa,D}(1 - l_A))$

System State: We represent the system state as the pair (i, j) , where i and j represent the current backoff stages of transmitters A and B , respectively. Note that $0 \leq (i, j) \leq m$, where $m + 1$ is equal to the maximum retransmission limit. The key approximation in our model is that, at each switching time, the next state does not depend on the current state. This allows us to model the evolution of our bidimensional state process with a discrete-time Markov chain embedded over continuous time at the time instants in which *both senders* can potentially start transmitting the first packet of a new data exchange. These packets could be RTS packets with the four-way handshake or data packets with basic access.

We further assume that a station's backoff counter is geometrically distributed over the contention window. This allows us to exploit the memoryless property of the geometric distribution without accounting for the remaining number of backoff slots. The geometrically distributed backoff counter at stage i is given by $\gamma_i = (2/W_i - 1)$, where W_i is the window size of backoff stage i . Consequently, a station in stage i attempts a new transmission with probability γ_i .

Transition Probability Calculation: Nodes A and B have transmission probabilities of γ_i and γ_j , corresponding to backoff stages i and j , respectively. The transition probabilities stem from the generic state (i, j) and are summarized in Tables I–IV for both types of access mechanisms and cross-flow connectivities (refer to Fig. 5).

Basic Access with Symmetric Cross-Flow Connectivity: We now calculate the transition probabilities with basic access (i.e., with RTS/CTS disabled) and symmetric cross-flow connectivity. We have summarized the transition probabilities for this group in Table I. The first row is an idle state in which neither of the nodes is transmitting a new packet. The second row refers to a single-access state leading to a successful transmission by A . The probability of this event is equal to the probability that: 1) sender A transmits and sender B does not attempt to transmit during A 's transmission, $\gamma_i(1 - \gamma_j)^{f_A}$, where f_A denotes A 's data packet size in mini-slots; and 2) A 's data packet is received successfully by its receiver, $1 - l_A$, where l_A is the loss probability based on the link quality. The corresponding unsuccessful single-access state happens with the probability $\gamma_i(1 - \gamma_j)^{f_A}l_A$ due to A 's packet loss. Similar single-access state probabilities can be calculated for B due to symmetry. On the other hand, if the state is neither an idle state nor a single-access state, then it would be an overlapping-packets state. We denote the probability that the system enters such a state by P_θ , which is equal to

$$P_\theta = 1 - (1 - \gamma_i)(1 - \gamma_j) - (\gamma_i(1 - \gamma_j)^{f_A} - (\gamma_j)(1 - \gamma_i)^{f_B}). \quad (2)$$

Since the later packet can arrive anywhere during the transmission of the first packet, the probability that it arrives during the synchronization bits of the first packet is negligible. Thus, with basic access, we assume that the later arriving data packet

TABLE II
USING THE FOUR-WAY HANDSHAKE WITH SYMMETRIC CONNECTIVITY

To State	Probability
$0, j+1$	$P_\theta S_{BA} O_{BA} C_{ABa} (1 - C_{BAb}) (1 - l_A)$
$0, j+k$	$P_\theta S_{BA} O_{BA} C_{ABa} C_{BAb} C_{Aba} C_{abA} (1 - C_{baB}) P_{k-1}(rts) C_{ABa,D}^{U(k-1)} (1 - l_A)$
$i+k, 0$	$P_\theta S_{BA} C_{BAb} (1 - O_{BA} C_{ABa} C_{Aba} C_{baB}) P_{k-1}(rts) C_{BAb,D}^{U(k-1)} (1 - l_B)$
$i+k, 0$	$P_\theta S_{BA} C_{BAb} O_{BA} C_{ABa} C_{Aba} C_{baB} (1 - C_{abA}) P_{k-1}(rts) C_{BAb,D}^{U(k-1)} (1 - l_B)$
$0, 0$	$P_\theta S_{BA} C_{BAb} O_{BA} C_{ABa} C_{Aba} C_{baB} C_{abA} C_{BAb,D} C_{ABa,D} (1 - l_B) (1 - l_A)$

is always dropped. The next two rows of Table I calculate the transmission probabilities when A 's packet arrives earlier than B 's packet. With B leading A , the transition probability can be calculated similarly due to symmetry.

Let $C_{XYx,D}$ denote the probability that X 's data packet is captured over Y 's transmission at receiver x . The third row of our table denotes a successful transmission by A when it overlaps with B 's transmission. This probability is equal to the probability that: 1) packets overlap, P_θ ; 2) A 's packet arrives earlier or exactly at the same time as B 's packet conditioned that the packets overlap, S_{AB} ; 3) A is captured, $C_{ABa,D}$; and 4) A 's packet is not lost due to poor channel conditions, $1 - l_A$. The probability that B 's packet arrives later than A 's packet conditioned that they overlap, S_{AB} , is equal to

$$\frac{\gamma_i \gamma_j \left(1 + \sum_{k=1}^{f_A-1} (1 - \gamma_j)^k \right)}{\gamma_i \gamma_j \left(1 + \sum_{k=1}^{f_A-1} (1 - \gamma_j)^k + \sum_{k=1}^{f_B-1} (1 - \gamma_i)^k \right)}. \quad (3)$$

The first term in the numerator of (3) is the probability that both packets are transmitted at the same time. The second term is the probability that B 's packet arrives later. The denominator is the overlapping-packets state probability.

A 's packet would be lost if poor channel conditions exist or the packet is not captured. This results in a backoff stage of $(i+1, j+1)$ and is calculated in the fourth row of Table I. Finally, in order to calculate the state duration for overlapping states, we assume that the later packet arrives in the middle of the first packet (which occurs on average). We select the overall length as the state duration.

Four-Way Handshake With Symmetric Cross-Flow Connectivity: We now calculate the transition probabilities when the RTS/CTS mechanism is enabled. The idle and single-access states can be calculated the same as for basic access by replacing f with f' , where f' is the transmission duration of an RTS in mini-slots. As a result, P_θ in (2) would correspond to the probability of overlapping RTS packets.

Fig. 6 depicts combinations in which B 's RTS arrives earlier than A 's RTS and at least one flow has a successful RTS/CTS exchange. In Cases 1 and 2, A has a winning RTS/CTS exchange. Cases 3 and 4 refer to a winning RTS/CTS exchange by B . In the last case, both senders receive successful CTS packets, and hence, both transmit data packets.

For each case in Fig. 6, the corresponding row in Table II calculates the transition probability. The transition probability of Case 1 is equal to the probability that: 1) RTS packets from A and B overlap, P_θ ; 2) B 's RTS arrives earlier conditioned that they overlap, S_{BA} , calculated by (3) with f' instead of f ; 3) A 's RTS arrives during the synchronization bits of B 's RTS, conditioned that they overlap and B leads A , O_{BA} ; 4) A 's RTS is captured at its receiver a over B 's RTS, C_{ABa} ; 5) B 's RTS is not captured at its receiver, $1 - C_{BAb}$; and 6) finally, the data




Fig. 6. Five cases for symmetric cross-flow connectivity based on timing and capture behaviors.

packet transmission is successful, $1 - l_A$, triggering an ACK transmission. Note that, in this case, the CTS packet transmitted by a is received by B , and hence, B defers for the rest of A 's transmission. Thus, the only remaining probability to calculate is O_{BA} , which equals

$$\frac{\gamma_i \gamma_j \sum_{k=0}^{s-1} (1 - \gamma_i)^k}{\gamma_i \gamma_j \sum_{k=0}^{f_B-1} (1 - \gamma_i)^k}. \quad (4)$$

The numerator of (4) is the probability that packets overlap and the later packet arrives during the synchronization bits of the first packet, where s is the duration of synchronization bits in mini-slots. The denominator calculates the probability that the later packet is transmitted anytime during the first packet's transmission. In Case 2 of Fig. 6, A 's CTS is retransmitted by b , but it is not received by B . Thus, sender B can retransmit RTS packets after a timeout. These RTS packets will not be captured at their receiver since the other flow is transmitting a data packet, and RTS packets arrive in the middle of its transmission. If $k-1$ further RTS packets are transmitted by B , the final backoff stage of j at the end of the transmission of the other flow will increase by k . The second row of Table II calculates the transition probability for Case 2. For a successful transmission by A , the data packet should be captured over the RTS retransmission, $C_{ABa,D}^{U(k-1)}$, where $U(k-1)$ is the unity function and is equal to one if any retransmissions happen, and zero otherwise. $P_n(rts)$ calculates the probability of n additional RTS packet transmissions and is equal to (5), shown at the bottom of the next page.

In (5), L is equal to the number of available transmission opportunities in mini-slots. For example, in Case 2 of Fig. 6, this duration is equal to the duration of the data and ACK exchange of flow Aa minus the timeout duration. RTS retransmissions only occur after a fixed timeout, and r is equal to the RTS plus timeout duration in mini-slots.

Equation (5) calculates the retransmission probability by adding all combinations in which n packets of size r can fit in L mini-slots. This expression divides the whole duration into n parts, each of size $m_i + r$ mini-slots, where B transmits after the first m_i slots. Since the retransmitted RTS packets fail, the backoff stage of the transmitter increases and will be reset to the minimum backoff once the maximum is reached.

Also relevant to Case 2 of Fig. 6, if B is able to receive the synchronization bits of b 's CTS packet but the MAC frame is lost, then B must wait an Extended Interframe Space (EIFS) before it restarts contending for the medium [16]. The EIFS duration is defined by the IEEE 802.11 standard to deal with incorrect MAC frame reception. According to the standard, if the PHY is able to indicate to the MAC that a frame transmission began but did not result in the correct reception of a complete MAC frame, then the receiving MAC should wait an EIFS period. The EIFS provides sufficient time for another station to acknowledge the incorrectly received frame. In Fig. 6, we have identified the cases in which EIFS might be used by marking the incorrectly received packet with a cross.

The remaining rows in Table II calculate the transition probabilities for the corresponding cases. These transition probabilities account for successful transmission by the flow with a successful RTS/CTS exchange. Hence, there must be a corresponding state that accounts for unsuccessful data transmissions which can be calculated from their corresponding successful transmissions. If A 's RTS is leading, the overlapping-packets state probabilities can be easily calculated from Table II due to topological symmetry. Finally, a state should account for overlapping RTS transmissions where neither transmitter receives a successful CTS packet. The resulting backoff stage would be $(i + 1, j + 1)$, and its probability is one minus the summation of all other states.

Basic Access With Asymmetric Cross-Flow Connectivity: We now move to the topology with asymmetric cross-flow connectivity and basic access. In this topology, a will not receive B 's transmissions, whereas b receives transmissions by A and a [refer to Fig. 5(b)]. The overlapping-packets state probabilities are summarized in Table III, and the nonoverlapping states remain the same as the symmetric scenario.

When packets overlap, a will only receive the packet sent by A , whereas b receives data packets transmitted by both A and B . As a result, if A 's packet arrives earlier and is not lost due to channel conditions, it will be successfully received at its receiver (first row of Table III). If B 's packet arrives earlier, three different scenarios can happen: simultaneous successful

TABLE III
BASIC ACCESS WITH ASYMMETRIC CONNECTIVITY

To State	Probability
$0, j + 1$	$P_\theta S_{AB}(1 - l_A)$
$0, 0$	$P_\theta S_{BA} C_{BAb,D} C_{Aba,D} (1 - l_A)(1 - l_B)$
$0, j + 1$	$P_\theta S_{BA} (1 - C_{BAb,D} + C_{BAb,D} l_B) (1 - l_A)$
$i + 1, 0$	$P_\theta S_{BA} C_{BAb,D} (1 - C_{Aba,D} + C_{Aba,D} l_A) (1 - l_B)$

transmissions (second row), successful transmission only by A (third row), or successful transmission only by B (fourth row). Any other overlapping-packets states will result in a backoff stage increase by both flows, where the probability is equal to one minus the summation of all other probabilities. Finally, we assume that with overlapping packets the later packet arrives in the middle of the other flow's transmission with the overall length as the state duration for each case.

Four-Way Handshake With Asymmetric Cross-Flow Connectivity: The main transition probabilities of this group are summarized in Table IV, and a sample of timeline graphs are plotted in Fig. 7. The first three rows of the table correspond to Case 1, in which A has a successful data transmission while the other flow retransmits RTS packets. This can happen in a single-access state or overlapping-packets state with either node transmitting earlier.

The fourth row of Table IV and Case 2 of Fig. 7 correspond to the probability that B has a successful data packet transmission and arrives earlier. Furthermore, the CTS packet transmitted by b is captured over A 's RTS at a . As a result, future RTS transmissions by A will not be replied to by a since it is able to set its NAV timer correctly. Hence, the backoff stage of A will increase if it retransmits RTS packets. Note that the CTS packet should arrive during the synchronization bits of A 's RTS. Denoted by O'_{BA} , this probability can be derived from (2). On the other hand, if the CTS packet transmitted by B is not captured over A 's RTS at a , different states happen depending on additional attempts of A and on successful or failed transmissions of each flow. One such example is plotted in Case 3 and calculated in the fifth row, where we have presented the probability of an additional attempt which results in simultaneously successful transmissions. Other states include no additional attempts by A , loss by either flow, or loss by both flows. The sixth row of the table calculates the same probability of Case 3 except the CTS packet is unsuccessful.

In the last case of Fig. 7, node A transmits earlier and captures the CTS packet transmitted by b . As a result, no further attempt is made. This probability is calculated in the seventh

$$\begin{aligned}
 P(n \text{ retransmissions}) = & \sum_{m_1=0}^{L-(n-1)r-1} (1 - \gamma_j)^{m_1} (\gamma_j) \times \left[\sum_{m_2=0}^{L-(n-1)r-m_1-1} (1 - \gamma_{j+1})^{m_2} (\gamma_{j+1}) \right. \\
 & \times \dots \times \left[\sum_{m_n=0}^{L-(n-1)r-1-\sum_{i=0}^{n-1} m_i} (1 - \gamma_{j+n-1})^{m_n} (\gamma_{j+n-1}) \right. \\
 & \times (1 - \gamma_{j+n})^{\frac{(L-(n-1)r-1-\sum_{i=1}^n m_i) + (L-(n-1)r-1-\sum_{i=1}^n m_i)}{2}} \left. \right] \left. \right] \quad (5)
 \end{aligned}$$

TABLE IV
USING THE FOUR-WAY HANDSHAKE WITH ASYMMETRIC CONNECTIVITY

To State	Probability
$0, j+k$	$\gamma_i(1-\gamma_j)^{f_A}P_k(rts)(1-l_A)$
$0, j+k$	$P_\theta S_{BA}(1-C_{BAb})(1-l_A)$
$0, j+k$	$P_\theta S_{AB}(1-O_{AB}C_{BAb}C_{BAb})P_{k-1}(rts)(1-l_A)$
$i+k, 0$	$P_\theta S_{BA}C_{BAb}(1-C_{Aba})O'_{BA}C_{bAa}P_{k-1}(rts)C_{BAb,D}^{U(k-1)}(1-l_A)$
$0, 0$	$P_\theta S_{BA}C_{BAb}(1-C_{Aba})(1-O'_{BA}C_{bAa})(1-P_0(rts))C_{BAb,D}C_{BAb,D}(1-l_B)C_{Aba}(1-l_A)$
$0, 0$	$P_\theta S_{BA}C_{BAb}C_{Aba}(1-C_{abA})(1-P_0(rts))C_{BAb,D}C_{BAb,D}(1-l_B)C_{Aba}(1-l_A)$
$i+1, 0$	$P_\theta S_{AB}O_{AB}C_{BAb}C_{BAb}C_{baA}(1-l_B)$
$0, 0$	$P_\theta S_{AB}O_{AB}C_{BAb}C_{BAb}(1-C_{baA})(1-C_{baA})(1-P_0(rts))C_{BAb,D}C_{BAb,D}(1-l_B)C_{Aba}(1-l_A)$




Fig. 7. Four cases for asymmetric cross-flow connectivity based on timing and capture behaviors.

row of Table IV. If the CTS packet is not captured, different outcomes can happen similarly to Case 3. The last row of the table calculates the probability if A makes another attempt and is successful. With certain capture relationships, two flows can have simultaneously winning RTS/CTS transmissions and hence, data packet transmissions. These probabilities can be calculated by plotting the timeline graphs, which we have omitted due to space limitations.

We emphasize that all the probabilities presented in Table IV assume a successful transmission by the flow winning RTS/CTS, and unsuccessful transmissions can be derived from them. Finally, any other overlapping-packets states will result in backoff stage increases by both flows, where the probability is equal to one minus the summation of all other probabilities.

Throughput Calculation: By numerically solving the Markov chain for each access mechanism and topology, which is ergodic for any choice of parameters, we obtain the stationary distribution $\pi = \pi_i, \forall i$. Long-term performance metrics such as throughput can be obtained directly from the solution of the Markov chain. From renewal-reward theory, the throughput of either flow is given by

$$T = \frac{\sum_n \pi_n P_{S_n}}{\Delta} \quad (6)$$

Here, P_{S_n} is the probability of successful transmission of either flow at state n , and Δ is the average duration of a step. Δ is computed from the average duration of all possible events in all states, weighted by their respective probabilities.

Handling Nonsaturated Flows: So far in our analysis, we have assumed that when the backoff counter of a flow reaches zero, the transmitter always sends a data packet, i.e., the senders are fully backlogged. We now extend our analysis to the case that the packet arrival rate of each flow i is λ_i . We define a new probability ρ_i , which is the probability that the sender has

a data packet to send when it is attempting to transmit a packet. To do so, we replace γ_i in our prior equations with $\gamma_i \times \rho_i$. With saturated throughput, ρ_i is equal to one. With unsaturated throughput, the achieved throughput of a flow i is less than or equal to λ_i .

With coupled sources, a closed-form expression for ρ_i that yields throughput equal to λ_i does not exist. Hence, we approximate ρ_i^{new} as $\alpha \times \rho_i^{\text{old}} + (1 - \alpha) \times (\lambda_i / T_i^{\text{old}})$ and adopt a global iterative procedure to update it. During each iteration, we utilize the throughput analysis to update the variables of every node as a function of its neighbor's ρ values (as computed in the previous iteration). The procedure ends when the throughput achieved by each flow is less than or equal to its demand.

Extending the Model to Other Technologies: While we derived the transition probabilities of the joint channel state evolution assuming a single capture threshold (as in 802.11b), the model can be easily extended to predict the throughput for other standards. For example, the additional capture threshold introduced by an MIM implementation (see Section II-B) can be taken into account in Tables I–IV with modifications to the state transitions and corresponding transition probabilities with the same approach. Additionally, one can extend our model to incorporate other medium access modifications. For example, the specifications in [17] may not allow for instant access to the channel after a successful transmission by the same node. This can be included in our model by lowering the transition probability of the previously successful transmitter.

Modeling Sources With Uncoupled Backoff: If an uncoupled, interfering source is within the sensing range of flow i , the later sender will sense the other's transmission and defer, which can be predicted by prior models [14]. Hence, we now discuss a hidden, uncoupled source. Furthermore, as the backoff evolutions are uncoupled, we use a decoupling technique to model the behavior of flow i . Namely, we model the private channel view evolution of the sender of flow i as a renewal process with three different states: 1) idle channel; 2) channel occupied by a successful transmission of flow i ; and 3) channel occupied by packet collisions.

Since we initially focus on a single hidden source, the sender's busy time is zero. To characterize the impact of an uncoupled source, we use a similar bidimensional Markov chain and transition probability as Table I. However, we make the following changes. First, as an uncoupled sender's backoff state does not change due to the other flow, we set the backoff value of flow j equal to zero. Second, since the sender of an uncoupled, interfering source does not receive ACKs, this value should be subtracted when calculating the state duration.

Let p be the probability that the transmission of a station is not successful. Then, the occurrences of the channel states are the

following: $\Pi_\sigma = 1 - \gamma_i$, $\Pi_s = \gamma_i$, $(1 - p)$, and $\Pi_c = \gamma_i, p$. Using standard renewal-reward theory, the throughput of the node is given by

$$T_P = \frac{\gamma_i(1 - p)}{\Pi_s T_s + \Pi_c T_c + \Pi_\sigma \sigma}. \quad (7)$$

Now, the transmission attempt probability, γ_i , is a deterministic function of p given by [14] and is equal to

$$\frac{2q(1 - p^{m+1})}{q(1 - p^{m+1}) + W_0 [1 - p - p(2p)^{m'}(1 + p^{m-m'}q)]}. \quad (8)$$

Here, $q = 1 - 2p$, W_0 is the minimum window size, m is the *maximum retry limit*, and m' is the backoff stage at which the window size reaches its maximum value ($m' \leq m$). The average duration of successful transmission or collision can be computed *a priori* [14]. Thus, the only unknown variable in (8) is the conditional packet loss probability, p . If we assume that the hidden node is an ON-OFF process, p can be derived as follows. The ON period is equal to the packet transmission time that is fixed for a given modulation rate and packet size. The OFF period is an exponential random variable with an average duration of $\bar{T}_{\text{OFF}} = (1/\lambda_j)$, where λ_j is the packet arrival rate at the hidden node. We assume that the transmission attempts of i happen randomly in the ON-OFF process. Thus, p equals

$$1 - \frac{\bar{T}_{\text{OFF}}}{T_{\text{ON}} + \bar{T}_{\text{OFF}}} e^{-\frac{d}{\bar{T}_{\text{OFF}}}} - C_{ij} \frac{\bar{T}_{\text{OFF}}}{T_{\text{ON}} + \bar{T}_{\text{OFF}}} \left(1 - e^{-\frac{d}{\bar{T}_{\text{OFF}}}}\right). \quad (9)$$

Here, d is the duration of flow i 's packet, and successful transmissions occur: 1) when the first packet arrives and fits completely into the idle period of the on-off process; or 2) when the packet overlaps with another transmission and is captured. If the hidden node follows another transmission pattern, an appropriate value of p can be plugged in.

IV. URBAN EXPERIMENTAL EVALUATION OF INFORMATION AND CHANNEL ASYMMETRIES

In this section, we perform thousands of measurements of coupled flows in an urban mesh network to both validate our model and experimentally analyze the complex factors that contribute to different throughput-sharing modes. With our model, we explore the full set of interdependencies that lead to this behavior and show that reverse capture plays a critical role. Furthermore, we experimentally and analytically show that this inversion can be based on small-scale channel fluctuations common to urban environments.

A. Experimental Setup and Measured Model Inputs

Throughout the TFA experiments, we activate two fully backlogged UDP flows (Aa and Bb) with 1500-B packets. We repeat the experiment in 120-s intervals for all combinations of 802.11b rates and for both access mechanisms. Before the experiment, we measure the data packet loss probability per modulation rate for each flow in isolation for our model. During the throughput experiments, we perform per-second SNR measurements and use the average relative SNR per link pair. Our capture measurements from Section II are then used to find the corresponding capture probability. Tables V and VI describe the average relative SNR for each possible competing link pair for the topologies. In some cases, one of the two competing links lacks connectivity, which results in a capture probability of 1 for

TABLE V
SYMMETRIC CROSS-FLOW CONNECTIVITY SUBTOPOLOGY (POSITIVE VALUE FAVORS Aa , AND NEGATIVE Bb)

RX	Relative SNR (dB)			
	$A - B$	$A - b$	$a - B$	$a - b$
A	-	-	$+\infty$	+1.0
a	+0.6	-1.4	-	-
B	-	$-\infty$	-	+3.9
b	-3.2	-	+1.0	-

TABLE VI
ASYMMETRIC CROSS-FLOW CONNECTIVITY SUBTOPOLOGY (POSITIVE VALUE FAVORS Aa , AND NEGATIVE Bb)

RX	Relative SNR (dB)			
	$A - B$	$A - b$	$a - B$	$a - b$
A	-	-	$+\infty$	+3.8
a	-7.1	-2.8	-	-
B	-	$-\infty$	-	$-\infty$
b	-1.6	-	-2.6	-

the other link. We denote this as $+\infty$ dB or $-\infty$ dB in Tables V and VI.

B. Channel Asymmetries With Symmetric Connectivity

As a baseline for our model validation, we first consider the throughput of coupled flows with symmetric cross-flow connectivity, which has been shown to fairly share bandwidth in idealized channel conditions with equal modulation rates [4], [10]. While this topology has *symmetric* cross-flow connectivity, there is vast heterogeneity in channel conditions between the flows, resulting in diverse capture characteristics based on the packet size and modulation rate (see Section II). As an example, Table V shows the SNR matrix for the left topology of Fig. 5. Sender A sends to receiver a , and B sends to b . The SNR difference between A and B ($A - B$) is -3.2 dB at b , and $A - B$ is $+0.6$ dB at a . Hence, with overlapping control packets transmitted by A and B , the probability for B 's control packets to win capture at b is 0.98 from Fig. 3. The same packets from B are likely to collide with A 's packets at a since the channel differences are small and the resulting ability of A to capture is negligible. We now compare the measured throughput results to the predicted throughputs by our model and the model developed in [4]. Our model considers both channel and information asymmetries in order to predict the throughput sharing of the coupled flows, while the model developed in [4] does not consider channel asymmetries or the resulting diverse capture characteristics.

Fig. 8(a) and (b) depicts the throughput achieved by each flow with RTS/CTS enabled. We observe that our model, *Joint Asymmetries*, provides an excellent match with the measurements, while the model developed in [4], *Information Asymmetry*, cannot correctly predict the resulting throughput. The throughput for flow Aa is near zero for all rate combinations. Our model reveals the reason for the severely imbalanced throughput. For flow Aa to have a successful transmission, its RTS transmission should not overlap with RTS transmissions of B . Moreover, if B 's RTS arrives earlier, it will be captured while if A 's RTS overlaps with B , it will be dropped. As a result, A 's backoff stage will continuously increase, whereas B 's backoff stage remains close to zero. In Fig. 8(c) and (d), we perform similar analysis with RTS/CTS disabled to show that our model is accurate and find that flow Aa achieves nearly zero throughput in this case as well. With basic access,




Fig. 8. Inverted throughput-sharing modes from model and validated by urban experiments on the topology and channel conditions represented by Tables V and VI with symmetric and asymmetric cross-flow connectivity with basic access (two-way) and RTS/CTS (four-way). Here, E represents 11 Mb/s. (a) Node A: four-way symmetric. (b) Node B: four-way symmetric. (c) Node A: two-way symmetric. (d) Node B: two-way symmetric. (e) Node A: four-way asymmetric. (f) Node B: four-way asymmetric. (g) Node A: two-way asymmetric. (h) Node B: two-way asymmetric.

it is even more likely that A 's data packet overlaps with B 's data packet, resulting in capture for B at b and loss for A at a . In summary, for both access mechanisms with symmetric cross-flow connectivity, forward traffic capture causes a bi-modal shift in the throughput-sharing mode where the channel asymmetry overwhelms the symmetry in information at each transmitter. Later, we discuss the additional effects of capture in the reverse direction in this topology.

C. Channel Asymmetries With Asymmetric Connectivity

We now consider coupled flows that compete with asymmetric cross-flow connectivity. Under perfect channels, this case yields one flow starving due to lack of information [4], [10]. Similar to the table for the symmetric subtopology, Table VI describes the competing links in the considered subtopology. Recall that with asymmetric cross-flow connectivity $a - b$ at B is $-\infty$ dB. We repeat the same experiment with this grouping of nodes and channel configuration.

Fig. 8(e) and (f) reports the measured throughput results with RTS/CTS enabled along with the predicted throughput results by our model, *Joint Asymmetries*, and the model developed in [4], *Information Asymmetry*. Surprisingly, our measurements show that the flow without information (Bb) is able to achieve approximately the same throughput as the flow with information (Aa). Thus, the coupled flows have an inverted throughput-sharing mode from [4]. This can be explained by the much larger ability to capture at b for B and A 's inability to capture at its own receiver versus b . As opposed to basic access, there are many more dependencies that are required to allow equal sharing that we explore in the next section. In short, the joint presence of forward *and* reverse traffic is required to invert the imbalanced sharing of the topology, making it balanced. In Fig. 8(g) and (h), RTS/CTS is disabled. Here, we see that the throughput sharing is also approximately balanced. However, capture in the forward traffic direction alone is sufficient to balance the throughput with asymmetric cross-flow connectivity. *Finding: With asymmetric cross-flow connectivity, forward*

capture inverts the throughput-sharing mode for basic access. However, forward and reverse capture is required to invert the throughput-sharing mode for the four-way handshake.

D. Reverse Capture Shifts the Symmetric Profile

We now explore the full range of the aforementioned interdependencies to invert the throughput-sharing modes. To do so, we use our model and Jain's Fairness Index, defined as $(\sum x_i)^2 / (n \cdot \sum x_i^2)$, where x_i is the achieved throughput of flow i and n is the total number of flows [18]. The fairness index of 1 corresponds to an equal throughput sharing, whereas a fairness index of 0.5 corresponds to one flow starving and the other obtaining all the throughput.

Even when coupled flows have symmetric information, the throughput-sharing modes can be inverted for both types of access mechanisms. We now explore the impact that capture relationships have on the throughput sharing of the symmetric cross-flow connectivity topology. While we expect that the forward traffic capture would dominate the behavior of basic access, the role of reverse traffic on sharing is previously unstudied, especially with RTS/CTS enabled. Here, we present the results from our model where two coupled UDP, fully backlogged flows compete with a modulation rate of 5.5 Mb/s.

Fig. 9(a) depicts the fairness achieved by the two transmitters (A and B) based on their ability to capture at their own receivers (C_{BAb} and C_{ABa}) using basic access. The fairness property depends on the symmetry of the forward traffic capture of the two flows. Prior work has predicted three points of this figure: [2] predicted the left and right corners of Fig. 9(a) (starvation mode), and [4] and [10] predicted the result with no capture (fair-sharing mode).

With the four-way handshake, the fairness index remains nearly identical to Fig. 9(a). In fact, if the graphs for both access mechanisms were overlaid, the differences are almost indistinguishable. However, the capture requirements for each access mechanism is very different (i.e., for a given channel condition, RTS packets have much lower capture thresholds




Fig. 9. Symmetric cross-flow connectivity with (a) two-way symmetric access and (b) four-way symmetric access with both reverse captures of Bb to 1.

than data packets). *Finding:* For a given channel condition, use of two-way versus four-way handshake can yield a bimodal shift because of the increasing ability to perform physical layer capture by the lower modulation rate and smaller size of the RTS packet as compared to the data packet.

Fig. 9(d) depicts our model's throughput prediction for the symmetric case with the four-way handshake where the reverse traffic is fully able to capture (C_{bAa} and C_{baA} equal 1). We observe that a shift in the sharing occurs, favoring flow Bb (which is able to capture in the reverse direction). To achieve a balanced throughput in this case, flow Aa must have a greater forward traffic capture than Bb . *Finding:* Reverse traffic capture shifts the throughput-sharing mode with symmetric cross-flow connectivity and the four-way handshake.

E. Forward and Reverse Capture With Information Asymmetry

With asymmetric cross-flow connectivity, we showed that fairness occurred when the information-poor flow Bb (i.e., the flow which lacks information) is able to capture in the forward traffic direction (C_{BAb}) for basic access. However, we have not yet considered the effect on the sharing when the information-rich flow Aa also has forward traffic capture (C_{ABa}).

Fig. 10(a) depicts a three-dimensional diagram of the fairness of two transmitters, A and B , according to their respective ability to capture at receivers a and b . In the left part of the figure, the information-poor node (B) is able to completely capture at b and A is unable to capture at a . This is the scenario that leads to perfect sharing for basic access. As A 's forward traffic capture (C_{ABa}) increases, the fairness index decreases rapidly and independent of B 's forward traffic capture value (C_{BAb}). *Finding:* With asymmetric cross-flow connectivity and basic access, inversion of the throughput-sharing mode primarily depends on the information-rich flow losing forward traffic capture and secondarily depends on the information-poor flow winning forward traffic capture.

Now, we consider the asymmetric case with the four-way handshake. Here, all four directions of capture in the forward and reverse directions must be considered since the RTS/CTS exchange preempts any data transmission. For the information-poor flow (Bb), the most important relationships to equalize throughput sharing is the forward over forward traffic capture (C_{BAb}) and the reverse over forward traffic capture (C_{bAa}). We first present the results from the model with these two capture relationships. We later show other relationships that contribute to increased throughput of Bb .

Fig. 10(b) depicts the fairness index for asymmetric cross-flow connectivity with the four-way handshake based upon the ability of the information-poor flow Bb to capture in

the forward direction versus competing forward traffic (C_{BAb}) and in the reverse direction versus competing forward traffic (C_{bAa}). On the left and right corners of the figure, near starvation of flow Bb occurs with the complete capture of the reverse or forward direction, respectively. However, in the middle of the figure, both relationships winning capture contribute to a much more equivalent throughput sharing. In Fig. 10(c), we add the ability of Bb to capture in the forward direction versus the reverse traffic (C_{BAb}). In the middle of the figure, we observe that a completely fair distribution of throughput (and complete inversion) can now be experienced. Finally, in Fig. 10(d), we add the ability of Bb to capture in the reverse direction versus reverse traffic (C_{baA}). Where these four capture relationships are 1 (middle of the figure), we observe that the flow Bb actually achieves greater throughput than Aa . *Finding:* With asymmetric cross-flow connectivity and the four-way handshake, the information-poor flow requires a confluence of link capture relationships to cause the throughput-sharing mode to invert. Yet, when the throughput-sharing mode does invert, the information-poor transmitter can obtain even higher throughput than the information-rich transmitter, a behavior that does not occur with basic access.

F. Small-Scale Channel Fluctuations Driving Modes

From our thousands of measurements over the course of a month on multiple topologies, we found many topologies have highly varying throughput sharing. The vast differences are despite the use of off-peak times for our experiments and limited activity of other nodes in the mesh network. In a particular grouping of four nodes with asymmetric cross-flow connectivity (described in Table VI), we found that the throughput sharing over a month's time period fluctuated from a starvation mode to a fair-sharing mode.

In Fig. 11(a), RTS/CTS is enabled, and each flow's throughput is depicted over the course of a month where the average throughput is represented by a bar, and the standard deviation is represented by error bars above and below the average. We observe that across many different modulation rate combinations, both flows have highly varying achieved throughput. Fig. 11(b) shows single-day measurements from the same scenario where both the minimum and maximum difference between throughput sharing were achieved, each with a data rate of 11 Mb/s. With the minimum difference in sharing (Day 1), nearly equal throughput is achieved, while the case with the maximum difference (Day 2) is severely imbalanced. By examining the differences in average SNR values between the two experiments, we observed a 1-dB relative difference in the pair of competing links in the forward traffic direction with negligible differences for all other links. Namely, when flow Bb is able to win forward traffic capture, it can achieve approximately equal throughput with flow Aa (as described with both access mechanisms in the validation experiment above). However, when B is unable to do so (e.g., when $A - B$ at b in Table VI goes from -1.6 to -0.6 dB), there is a large throughput difference between flows Bb and Aa . Recall that reverse traffic capture is present in this topology, allowing the forward traffic capture relationship to make a difference. Therefore, small-scale channel fluctuations allow switching between the fair-sharing and starvation modes. While the change in channel state is between different days of the month (larger timescale), it exposes a corresponding




Fig. 10. Asymmetric cross-flow connectivity with (a) two-way handshake and (b)–(d) four-way handshake, where the information-poor flow has increased chance to win the contention due to favorable capture relationships.




Fig. 11. Asymmetric topology (from Table VI) over a month of measurements and measurements from the two days where the difference in throughput achieved is minimum (Day 1) and maximum (Day 2). (a) High Variation (over a month). (b) Throughput (on two different days).




Fig. 12. Symmetric topology (Table V) over a month of measurements and measurements from the two days where the difference in throughput achieved is minimum (Day 1) and maximum (Day 2). (a) High variation (over a month). (b) Throughput (on two different days).

throughput-sharing behavior that is caused by the small-scale fluctuations in the TFA network. *Finding: Even common, small-scale channel fluctuations can cause a change in mode to invert the throughput-sharing mode.*

We observe a similar behavior with the symmetric topology described in Table V. In Fig. 12(a), we show that this topology also had highly varying throughput-sharing properties. When the throughput is nearly equivalent (Day 1), both senders of the flows are able to achieve physical layer capture at the intended receiver. However, a 1-dB fade causes Aa to be unable to perform physical layer capture, and the two flows have very different achieved throughput (Day 2). We confirm with the model that changing this one parameter can drive such a relationship (i.e., it is not due to poor channel conditions on one run of the experiment). Namely, A 's signal to a increases by 1 dB (from 0.6 to 1.6) and causes the capture probability of the RTS message to go from 32% to 87%. Before the change and after the change, B kept a greater than 2 dB SNR difference over A at b , and the capture threshold of its RTS packets remained nearly 1. However, A being able to capture at its own receiver a with high probability allowed a nearly balanced throughput sharing. *Finding: The throughput imbalance induced by unequal forward traffic capture relationships with symmetric cross-flow connectivity disallows even the four-way handshake to establish balanced throughput sharing.* These findings show that an understanding of the capture behavior of the card and logging of SNR and achieved throughput over time (perhaps in off-peak hours) can reveal the reasons for throughput-sharing anomalies.

V. APPLICATIONS OF ANALYTICAL MODEL AND EXPERIMENTATION

In this section, we apply our experimental analysis and model in two ways. First, we consider how modulation rate can be selected according to joint properties of channel condition,

topology, and capture. Second, we predict and explain the disproportionate effect of low-rate control traffic on embedded data flows within a mesh network.

A. Capture-Induced Dimension to Rate Selection Problem

Prior work has considered the problem of choosing the modulation rate that achieves the highest throughput based on the channel condition from the sender to receiver [19]. However, no prior work has considered the interdependence of physical layer capture and modulation rate selection. We now apply our experimental analysis of different capture behaviors and our model to a scenario in which a data flow competes with an uncoupled, hidden source that saturates the channel. Using our analytical model, we fix the modulation rate (11 Mb/s) and packet size (1500 B) for the interfering transmitter. For the data flow, we vary the modulation rate (2, 5.5, and 11 Mb/s) and packet size (100, 500, 1000, and 1500 B) with fully backlogged, UDP traffic over a range of relative SNR.

Fig. 13 depicts the throughput for the data flow based upon its choice of modulation rate, packet size, and relative SNR to the interfering flow. At low relative SNR, no throughput is achieved as all packets are unable to capture against the interfering transmitter. As the relative SNR increases, the throughput increases based upon the capture probability. Observe that the 5.5-Mb/s rate is able to achieve the highest throughput out of all modulation rates for a relative SNR of at least 6 dB. This contrasts the throughput-maximizing modulation rate whenever the interfering source is off (11 Mb/s), based on the link quality from sender to receiver. Furthermore, it contrasts the throughput-maximizing modulation rate if the interfering packet was always on (i.e., the difference between the data flow's packet and the interfering packet becomes the new channel condition), which would be 1 Mb/s. The difference between both of these cases is that capture allows




Fig. 13. Achievable throughput of a data flow when a hidden source with uncoupled backoff saturates the channel.

a temporarily “clean” channel until the receiver locks on, and a “noisy” channel (up to the capture threshold) thereafter. Thus, the throughput-maximizing modulation rate is not 11 nor 1 Mb/s, but 5.5 Mb/s due to its increased capture probability (from 11 Mb/s) and increased rate (from 1 Mb/s).

While this is a relatively small region of the graph, consider the relative capture thresholds for 802.11a as shown in [7]. Most of the delivery ratio curves for different modulation rates are completely orthogonal, meaning that the delivery ratio goes from 0 to 1 over orthogonal SNR regions. Thus, the rate selection problem with 802.11a would be based heavily upon this capture-induced dimension when competing with other transmitters as opposed to the channel condition between sender and receiver alone. *Finding: With interfering flows, the throughput-maximizing modulation rate may be lower than the throughput-maximizing modulation rate allowed for the flow in isolation, and higher than the throughput-maximizing modulation rate for a channel with constant noise.*

As a result of these findings, a rate adaptation protocol could leverage a modification to the physical layer implementation that is similar to the MIM implementation of 802.11a to inform higher layers of overlapping packets and resulting capture events. The receiver could leverage such information directly in SNR-based protocols, as the decision is made by the receiver and fed back to the transmitter. Instead of making a rate decision based on SNR alone, the decision could be based on SNR and forward capture relationships that are detected. The implementation for rate adaptation protocols where the decision is made by the transmitter would require the addition of a feedback loop from the receiver into the protocol.

B. Predicting Throughput Reduction for Diverse Traffic Types

We now apply our model and experimentation to a second domain in which hidden nodes transmit small-sized, low-rate control overhead, causing a disproportionally large effect on large-sized, high-rate data flows. In contrast to the previous section on rate selection where the goal was to adapt the traffic parameters to achieve the highest throughput, here we are assuming the rate and packet size choices to be already made (e.g., by a particular application and QoS requirements thereof) and examine how these characteristics might be affected by an interfering source. While prior work has shown the existence of such an effect in a mesh network [13], the reasons for the losses have not clearly been identified, nor have the implications been explored for other packet types and modulation rates often used throughout a network. Since this effect can exist for all traffic types, we term it the *throughput reduction factor*. The




Fig. 14. Range of throughput reduction (a) factors and (b) post-reduction throughput based upon the transmission rate of the data flow.

throughput reduction factor is the achieved throughput without the presence of interfering traffic minus the achieved throughput with interfering traffic over the injected traffic rate of the interfering sources.

We begin by understanding the throughput reduction factor within TFA for different modulation rates for the data flow. We perform the experiments at off-peak times and generate the interfering traffic profile presented in [13] of 100-B packets at a rate of 10 kb/s from a hidden interferer. We measure the throughput reduction factor for data flows for each of the modulation rates for flows across the TFA network.

Fig. 14 depicts (a) the throughput reduction factor and (b) corresponding post-reduction throughput for each relative SNR from -5 to $+10$ dB over each modulation rate. In the left of Fig. 14(a), the data flow loses capture to the interfering source, experiencing the highest throughput reduction factor. In the right of Fig. 14(a), the data flow has minimal reduction due to higher capture probability at the receiver. Clearly, the higher modulation rates have greater penalties as the interfering source collides and forces the data flow to back off. However, the throughput reduction factor is also affected by the ability of the modulation rate to capture. For example, 1 Mb/s reaches a near-zero value at a relative SNR of 5 dB versus 8 dB for 11 Mb/s. Coupled with channel quality increases, the higher-order modulations can lose more than 2 Mb/s for the same competing tens of kilobits of offered load [from left to right of Fig. 14(b)]. *Finding: For a given relative SNR in relation to the interfering traffic, the modulation rate choice of the data flow can change a throughput reduction factor by over 30, resulting in more than 2 Mb/s of loss.*

We now use our measurements and model to reveal the key system properties that yield the throughput reduction factor. In our capture experiments, neither modulation rate nor packet size of the interfering source affected the performance of the data flow, i.e., if the data flow was stronger, the capture threshold did not depend on the traffic characteristics of the interfering traffic (given that it was overlapping). Thus, for the throughput reduction factor, neither packet size nor modulation rate of the interfering flow determines the capture behavior of overlapping data packets.² Hence, the throughput reduction factor is primarily driven by the traffic characteristics of the data flow since choices of modulation rate and packet size affect the ability to capture. Therefore, we now use our model to predict the throughput reduction factor for different packet sizes and modulation rates of the data flow with the same interfering traffic profile as before.

²The interferer’s offered load would affect the throughput reduction due to the increased probability of overlapping packets as the load increased.




Fig. 15. Throughput reduction factor for a data flow based on its packet size and physical layer rate against a low-rate interfering source.

Fig. 15 depicts the throughput reduction factor for the data flow based on its packet size and modulation rate over different relative SNR values versus the interfering source. The highest values of the throughput reduction factor are for large-sized, high-rate packets, and lowest for small-sized, low-rate packets. While this is somewhat expected due to the inability of small-sized packets to achieve high throughput, the crossover point of when the factor approaches zero is interesting. For example, consider a modulation rate of 2 Mb/s. For the packet size of 100 B, the throughput reduction factor goes to a near-zero value at +2 dB versus +7 dB for 1500-B packets at the same rate. *Finding: Our model and analysis shows that the largest throughput reductions are due to the joint factors of control traffic originating from a hidden terminal and the data traffic's inability to win capture over the control traffic. Note that since the control source has an uncoupled backoff with broadcast traffic, the control packet's ability to win capture does not affect the throughput of the data flow as with the interfering source with coupled backoff behavior.*

In this paper, we directly show how our model can be used to aid rate selection and predict multiplicative throughput reduction, which can lead to improved link capacity network-wide and a characterization of the net benefit of any protocol that induces overhead, respectively. However, since our model can accurately characterize throughput-sharing behavior, the model could also lead to fundamental advances for all types of protocols and algorithms including routing, network planning, and other forms of network management.

VI. RELATED WORK

CSMA Models: There is a rich body of work on modeling CSMA, dating back to the seminal work by Kleinrock and Tobagi [20]. With the introduction of IEEE 802.11, Bianchi presented a simplified model that used the assumptions of single-rate, single-clique, and fully backlogged traffic with fixed packet size [14]. More general topologies and scenarios were explored in [4] and [5], in which competing transmitters could have different views of the channel. The fully backlogged assumption was removed though the system was still single-rate with fixed packet size and binary interference. Furthermore, [21] considered physical layer features such as hidden terminals and capture without topological asymmetries. The modeled scenarios have been further generalized to include nonbinary interference with physical-layer capture characteristics [22] and interference from an arbitrary number of transmitters [23] for flows using the two-way handshake. In contrast, we consider the effects of heterogeneous modulation

rates and the resulting MAC and PHY layer properties that jointly impact the throughput sharing of competing flows.

Measurement-Based Models: The aforementioned works of [22]–[24] further distinguish themselves from prior works in that the models are measurement-based to balance the general applicability of modeling with the realism of actual networks, requiring only $O(n^2)$ measurements. In contrast, while they consider the performance of flows for the two-way handshake in random topologies, we show via measurement and modeling that the performance of both access mechanisms have greatly increased complexities due to symmetric and asymmetric connectivity between flows as well as heterogeneous traffic capture directionalities.

Measurement Studies of Multihop 802.11 Networks: A number of works have identified the channel conditions and timing under which physical layer capture occurs for pairs of nodes [1], [7], [8]. Others have proposed modifying physical layer properties to address the lack of fairness that results [27], [28]. Measurements have been performed on indoor multihop wireless topologies to characterize interactions of flows [9]. Finally, measurement studies have been performed in mesh networks to explore the link behavior [29], flow performance [30], rate adaptation [31], and overhead effects [13].

In contrast, while we identify where physical layer capture and topological disparities occur in our urban network, our focus is on how both jointly affect the throughput sharing of competing flows. Specifically, our work is the first to jointly study these information and channel asymmetries via modeling and experimentation to reveal that small-scale channel fluctuations common to an urban mesh deployment can yield bimodal performance shifts.

VII. CONCLUSION

In summary, we perform extensive measurements on coupled flows within an urban mesh network and analytically model the complex factors that exist with information and channel asymmetries. Our experimental analysis and model reveal that small-scale channel fluctuations common to urban environments are able to completely invert the throughput-sharing mode. Using our model, we explore the interdependencies of these complex factors and find that reverse capture plays a critical role in defining these performance modes. Lastly, we show how to extend and apply our model and experimentation to two different problem domains: modulation rate selection and the interaction of control and data traffic. There are immediate implications of our work. First, since the topological asymmetries (including the competing traffic flows' directionalities for capture relationships that form) greatly impact the most appropriate transmit parameter choices, networks would benefit from physical layer implementations that detect and inform higher layers about the capture events (i.e., a slight enhancement to the MIM implementation in 802.11a) to allow more informed adaptation decisions. Second, since the traffic parameter choices alter both capture relationships and information asymmetries, awareness of such a behavior enables prediction of new topological asymmetries that form as transmission parameters change, allowing throughput prediction for a broad array of future CSMA-based protocols and algorithms. As part of our future work, we plan to investigate the throughput-sharing behaviors with an additional number of competing flows.

REFERENCES

- [1] A. Kochut, A. Vasan, A. U. Shankar, and A. Agrawala, "Sniffing out the correct physical layer capture model in 802.11b," in *Proc. IEEE ICNP*, Oct. 2004, pp. 252–261.
- [2] F. Tobagi, "Modeling and performance analysis of multihop packet radio networks," *Proc. IEEE*, vol. 75, no. 1, pp. 135–155, Jan. 1987.
- [3] T. Nandagopal, T. Kim, X. Gao, and V. Bharghavan, "Achieving MAC layer fairness in wireless packet networks," in *Proc. ACM MobiCom*, Aug. 2000, pp. 87–98.
- [4] M. Garetto, J. Shi, and E. Knightly, "Modeling media access in embedded two-flow topologies of multi-hop wireless networks," in *Proc. ACM MobiCom*, Aug. 2005, pp. 200–214.
- [5] M. Garetto, T. Salonidis, and E. Knightly, "Modeling per-flow throughput and capturing starvation in CSMA multi-hop wireless networks," *IEEE/ACM Trans. Netw.*, vol. 16, no. 4, pp. 864–877, Aug. 2008.
- [6] C. Ware, J. Judge, J. Chicharo, and E. Dutkiewicz, "Unfairness and capture behaviour in 802.11 adhoc networks," in *Proc. IEEE ICC*, Jun. 2000, vol. 1, pp. 159–163.
- [7] J. Lee, W. Kim, S.-J. Lee, D. Jo, J. Ryu, T. Kwon, and Y. Choi, "An experimental study on the capture effect in 802.11a networks," in *Proc. ACM WinTECH*, Sep. 2007, pp. 19–26.
- [8] G. Judd and P. Steenkiste, "Understanding link-level 802.11 behavior: Replacing convention with measurement," in *Proc. WICON*, Oct. 2007, Article no. 19.
- [9] J. Lee, S.-J. Lee, W. Kim, D. Jo, T. Kwon, and Y. Choi, "Understanding interference and carrier sensing in wireless mesh networks," *IEEE Commun. Mag.*, vol. 47, no. 7, pp. 102–109, Jul. 2009.
- [10] V. Bharghavan, S. Demers, S. Shenker, and L. Zhang, "MACAW: A media access protocol for wireless LANs," in *Proc. ACM SIGCOMM*, Aug. 1994, pp. 212–225.
- [11] J. Camp, J. Robinson, C. Steger, and E. Knightly, "Measurement driven deployment of a two-tier urban mesh access network," in *Proc. ACM MobiSys*, Uppsala, Sweden, Jun. 2006, pp. 96–109.
- [12] R. Sinha, C. Papadopoulos, and J. Heidemann, "Internet packet size distributions: Some observations," University of Southern California, Los Angeles, CA, Tech. Rep. ISI-TR-2007-643, May 2007.
- [13] J. Camp, V. Mancuso, O. Gurewitz, and E. Knightly, "A measurement study of multiplicative overhead effects in wireless networks," in *Proc. IEEE INFOCOM*, Apr. 2008, pp. 76–80.
- [14] G. Bianchi, "Performance analysis of the IEEE 802.11 distributed coordination function," *IEEE J. Sel. Areas Commun.*, vol. 18, no. 3, pp. 535–547, Mar. 2000.
- [15] T. Rappaport, *Wireless Communications, Principles & Practice*, ser. Emerging Technologies, T. Rappaport, Ed. Upper Saddle River, NJ: Prentice-Hall, 1996.
- [16] *Wireless LAN Medium Access Control and Physical Layer Specification*, ANSI/IEEE Standard 802.11, IEEE, 2003.
- [17] *Standard 802.11e, Part 11 Amendment: Medium Access Method Quality of Service Enhancements*, ANSI/IEEE Standard 802.11, IEEE, 2005.
- [18] R. Jain, W. Hawe, and D. Chiu, "A quantitative measure of fairness and discrimination for resource allocation in shared computer systems," Ohio State University, Columbus, OH, Tech. Rep. DEC-TR-301, 1984.
- [19] G. Judd, X. Wang, and P. Steenkiste, "Efficient channel-aware rate adaptation in dynamic environments," in *Proc. ACM MobiSys*, 2008, pp. 118–131.
- [20] L. Kleinrock and F. A. Tobagi, "Packet switching in radio channels: Part I—Carrier sense multiple-access modes and their throughput-delay characteristics," *IEEE Trans. Commun.*, vol. COM-23, no. 12, pp. 1400–1416, Dec. 1975.
- [21] M. Carvalho and J. J. Garcia-Luna-Aceves, "A scalable model for channel access protocols in multihop ad hoc networks," in *Proc. ACM MobiCom*, Sep. 2004, pp. 330–344.
- [22] C. Reis, R. Mahajan, M. Rodrig, D. Wetherall, and J. Zahorjan, "Measurement-based models of delivery and interference in static wireless networks," in *Proc. ACM SIGCOMM*, Sep. 2006, pp. 51–62.
- [23] L. Qiu, Y. Zhang, F. Wang, M.-K. Han, and R. Mahajan, "A general model of wireless interference," in *Proc. ACM MobiCom*, Sep. 2007, pp. 171–182.
- [24] Y. Li, L. Qiu, Y. Zhang, R. Mahajan, and E. Rozner, "Predictable performance optimization for wireless networks," in *Proc. ACM SIGCOMM*, Aug. 2008, pp. 413–426.
- [25] B. Sadeghi, V. Kanodia, A. Sabharwal, and E. Knightly, "Opportunistic media access for multirate ad hoc networks," in *Proc. ACM MobiCom*, Sep. 2002, pp. 24–35.
- [26] G. Cantieni, Q. Ni, C. Barakat, and T. Turletti, "Performance analysis under finite load and improvements for multirate 802.11," *Comput. Commun.*, vol. 28, no. 10, pp. 1095–1109, 2005.
- [27] T.-Y. Lin and J. Hou, "Interplay of spatial reuse and SINR-determined data rates in CSMA/CA-based, multi-hop, multi-rate wireless networks," in *Proc. IEEE INFOCOM*, May 2007, pp. 803–811.
- [28] S. Ganu, K. Ramachandran, M. Gruteser, I. Seskar, and J. Deng, "Methods for restoring MAC layer fairness in IEEE 802.11 networks with physical layer capture," in *Proc. ACM REALMAN*, May 2006, pp. 7–14.
- [29] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris, "Link-level measurements from an 802.11 mesh network," in *Proc. ACM SIGCOMM*, Sep. 2004, pp. 121–132.
- [30] J. Bicket, S. Biswas, D. Aguayo, and R. Morris, "Architecture and evaluation of the MIT Roofnet mesh network," in *Proc. ACM MobiCom*, Aug. 2005, pp. 31–42.
- [31] J. C. Bicket, "Bit-rate selection in wireless networks," M.S. thesis, MIT, Cambridge, MA, Feb. 2005.



Joseph Camp (S'03–M'09) received the B.S. degree (with honors) from the University of Texas at Austin in 2003, and the M.S. and Ph.D. degrees from Rice University, Houston, TX, in 2005 and 2009, all in electrical and computer engineering.

He is an Assistant Professor of electrical engineering with Southern Methodist University (SMU), Dallas, TX. He joined the SMU faculty in 2009 after receiving the Ph.D. degree. From 2004 to 2009, he was the Chief Network Architect for the Technology For All (TFA) Network, a mesh deployment in Houston, TX, which serves thousands of users in an underresourced community. His research interests are in the areas of wireless networking and embedded systems, specifically focused on deployment, measurement, and analysis of large-scale wireless networks and development of embedded protocols for network hardware.

Dr. Camp received the Ralph Budd Award for the best engineering thesis at Rice University in 2010.



Ehsan Aryafar (S'05–M'11) received the B.S. degree in electrical engineering from the Sharif University of Technology, Tehran, Iran, in 2005, and the M.S. and Ph.D. degrees in electrical and computer engineering from Rice University, Houston, TX, in 2007 and 2011, respectively.

He is a Post-Doctoral Research Associate with Princeton University, Princeton, NJ. His research interests are in the areas of wireless networks, high-performance MAC protocol design, network deployment and resource provisioning, and network measurements. His current projects include white space networking and hybrid satellite and terrestrial networks.



Edward Knightly (S'91–M'96–SM'04–F'09) received the B.S. degree from Auburn University, Auburn, AL, in 1991, and the M.S. and Ph.D. degrees from the University of California, Berkeley, in 1992 and 1996, respectively, all in electrical engineering.

He is a Professor of electrical and computer engineering with Rice University. His research interests are in the areas of mobile and wireless networks and high-performance and denial-of-service resilient protocol design.

Prof. Knightly is a Sloan Fellow. He served as Associate Editor of numerous journals and special issues, including the IEEE/ACM TRANSACTIONS ON NETWORKING and IEEE JOURNAL ON SELECTED AREAS OF COMMUNICATIONS Special Issue on Multi-Hop Wireless Mesh Networks. He served as Technical Co-Chair of IEEE INFOCOM 2005, General Chair of ACM MobiHoc 2009 and ACM MobiSys 2007, and served on the program committee for numerous networking conferences including ICNP, INFOCOM, MobiCom, and SIGMETRICS. He is a recipient of the National Science Foundation CAREER Award. He received the Best Paper Award from ACM MobiCom 2008.